Causal Discovery as a Game

Frederick Eberhardt EBERHARDT@WUSTL.EDU
Department of Philosophy
Washington University in St. Louis
St. Louis, MO 63130, USA

Editor: Isabelle Guyon, Dominik Janzing and Bernhard Schölkopf

Abstract

This paper presents a game theoretic approach to causal discovery. The problem of causal discovery is framed as a game of the Scientist against Nature, in which Nature attempts to hide its secrets for as long as possible, and the Scientist makes her best effort at discovery while minimizing cost. This approach provides a very general framework for the assessment of different search procedures and a principled way of modeling the effect of choices between different experiments.

Keywords: causal discovery, interventions, search strategy, game theory, worst and expected case analysis

1. Introduction

In machine learning much of the literature on causal discovery has focused on discovery in passive observational data. The analysis of experimental data has been left to the field of experimental design, but there the focus has been on the optimal allocation of samples to a pre-determined set of treatment variables, and the subsequent analysis of the data. Very little work has been done on the selection of experiments. The specification of the best sequence of experiments to discover particular causal relations has largely been left to the "good judgment of the scientist." Only recently have first steps been taken to automate this process: Tong and Koller (2001); Murphy (2001); Yoo and Cooper (2003); Meganck et al. (2005) and He and Geng (2008) have presented approaches to select the next best experiment based on information theoretic measures or expected utility, and Eberhardt (2007) provided worst case bounds for such search strategies under different assumptions. In this paper a game theoretic analysis of sequences of experiments is proposed that identifies appropriate guidelines for the choice and comparison of different experimental strategies.

Randomized controlled trials (RCTs) are perhaps the most widely accepted standard to determine cause and effect. If, as intended by the randomization, the intervention makes the intervened variable independent of its normal causes, then it breaks any confounding of the causal effect of the intervened variable on the outcome variable by measured or unmeasured common causes.[1] Given a set of, say, three variables X, Y and Z, a scientist has many choices of

1. With regard to causal discovery, weaker forms of interventions can also provide insights but we will leave that issue aside here.

which variable(s) to randomize. She could intervene on any one and measure the other two. She could randomize any two, independently or not, and measure the third, etc. Whichever choice she makes, one experiment will in general not guarantee – even in the large sample limit – the discovery of the *true* causal structure among the 25 *possible* (directed acyclic) causal structures over the three variables. *Sequences* of different experiments are often necessary to determine all the causal relations between variables. But what is the best sequence, and in what sense of "best"?

2. Worst Case Analysis

One way to compare different search strategies is to consider their worst case performance. In Eberhardt et al. (2005) we gave worst case analyses of different search procedures for causal discovery involving different types of interventions under a variety of different assumptions. The quality of different search procedures was measured in terms of the number of experiments *sufficient and in the worst case necessary* to discover the true causal structure among N variables. The worst case was characterized by the causal structure that required the longest sequence of experiments that could not be avoided (by more appropriate choices of experiments given the available knowledge at the choice point). The following table summarizes the results for sequences of experiments with single or multiple simultaneous RCT-type interventions per experiment on a set of N causal variables, with and without latent variables:[2]

Interventions per Experiment	Latent Variables Present	Number of Experiments
Single	No	2 if $N = 2$ & $N-1$ if $N > 2$
Single	Yes	impossible
Multiple	No	$\lfloor \log_2(N) + 1 \rfloor$
Multiple	Yes	N

A worst case analysis provides an upper bound, but in practice the worst case may be very rare whereas a "typical" search problem might be resolved much faster. Consequently, the expected performance is often considered. The computation of an expectation depends on a distribution over the possible hypotheses. In the case of three variables, it would require a distribution over the 25 possible (acyclic) causal structures. In many cases, the uniform distribution is used, but without more specific knowledge of the domain under consideration, it is not clear why the uniform distribution is more appropriate than any other. Often sparsity assumptions play a crucial role in restricting the hypothesis space and it is not clear that a uniform distribution over hypotheses is "uninformative" when a sequence of experiments is used for discovery. What, then, can be said about an expected case performance without commitment to a particular distribution?

3. Expectation and Optimization

One approach supported by a game-theoretic interpretation of the discovery problem is the *worst case expected* performance, i.e. the upper bound on the expected length of sequences of experiments sufficient and in the worst case necessary to discover the causal structure, *no matter what the probability distribution over the set of directed acyclic graphs is*. That is, for

2. The second row indicates that no sequence of experiments, in which only a single variable is subject to an RCT-type intervention, is sufficient to discover the causal structure *in the worst case* if there are latent variables. Under different assumptions, such as linearity, discovery is possible (see Eberhardt (2007)).

each distribution $P(\mathcal{G})$ over the set $\mathcal{G}$ of directed acyclic graphs, take the expectation $E_P(.)$ of the number of experiments $\#ex(.)$ sufficient and in the worst case necessary to uniquely discover the true causal graph G, whatever G is. Then take the upper bound – the supremum – of those expectations. Or formally:

$$\sup_E E_P(\#ex(G)) \quad \text{over all} \quad P(\mathcal{G}). \tag{1}$$

The key to determining this quantity is the specification of $\#ex(G)$ for some true underlying causal structure G. To specify this quantity we need to specify how experiments are chosen. But how and which experiments are chosen affects which causal structures are difficult to learn, so the supremum is affected by both the underlying distribution over causal structures and the sequence of experiments that is used to identify which one is true.[3]

Let $\mathcal{S}$ be a strategy that specifies a sequence of experiments, in which the next experiment is determined with probability 1 contingent on the evidence revealed in all previous experiments. Given a set $\mathcal{G}' \subseteq \mathcal{G}$ of possible causal structures (determined by non-zero probability in the probability distribution $P(.)$ over causal structures), let $\#ex_{\mathcal{S}}(G)$ be the number of experiments according to strategy $\mathcal{S}$ that is necessary and sufficient to uniquely identify a particular causal structure $G \in \mathcal{G}'$. Since we are interested in an optimal number of experiments, we can now define $\#ex(G)$ as the number of experiments necessary and sufficient to uniquely identify $G \in \mathcal{G}'$ using a strategy $\mathcal{S}^+$ where

$$\forall \mathcal{S} \neq \mathcal{S}^+ \qquad E_P(\#ex_{\mathcal{S}}(G)) \geq E_P(\#ex_{\mathcal{S}^+}(G)). \tag{2}$$

That is, for a given set of possible causal structures $\mathcal{G}'$, $\#ex(G)$ specifies the expected number of experiments necessary and sufficient to uniquely identify the causal structure $G \in \mathcal{G}'$ using a most efficient strategy $\mathcal{S}^+$. However, since we are interested in the *supremum* of the expectations, the distribution $P(.)$ that specifies the set of possible graphs, must be such that it is implies the largest expected number of experiments for a given strategy, i.e. given a strategy S, $P^+(.)$ is chosen such that for any

$$P(.) \neq P^+(.) \qquad E_{P^+}(\#ex_{\mathcal{S}}(G)) \geq E_P(\#ex_{\mathcal{S}}(G)). \tag{3}$$

Definitions (2) & (3) make the interdependence between a search strategy and the distribution over hypotheses explicit: Given a hypothesis space one can specify the optimal search strategy. Given a search strategy one can specify the hypothesis space that will make search most difficult. We can thus rephrase the supremum in (1) above as the following optimization:

$$E_{P^*}(\#ex_{\mathcal{S}^*}(G)) \tag{4}$$

where it is simultaneously the case that for any

$$P(.) \neq P^*(.) \qquad E_{P^*}(\#ex_{\mathcal{S}^*}(G)) \geq E_P(\#ex_{\mathcal{S}^*}(G)) \tag{5}$$

and given the set of possible causal structures implied by $P^*(.)$,

$$\forall \mathcal{S} \neq \mathcal{S}^* \qquad E_{P^*}(\#ex_{\mathcal{S}}(G)) \geq E_{P^*}(\#ex_{\mathcal{S}^*}(G)). \tag{6}$$

3. For example: If one always intervenes on X first, then causal structures in which X is an effect (but not a cause!) of the other variables, are more difficult to discover because any incoming causal influence on X is destroyed by the intervention, and so the structure cannot be distinguished from one in which X is causally independent of the other variables. Consequently, a distribution that puts more weight on those graphs will be a candidate for the maximum expectation. But such a distribution results in a much lower expectation if the first intervention always intervenes on one of the causes of X (say Y), since the $Y \to X$ edge is discovered immediately.

For fixed sequences of experiments that specify a particular experiment for a given history of evidence there is no solution to the above double optimization. That is, for any specific strategy $\mathcal{S}$ there is a probability distribution $P(.)$ that maximizes the number of experiments with respect to $\mathcal{S}$ (in fact, the expectation can always be forced to the absolute worst case bound). However, an alternative strategy $\mathcal{S}'$ would do better on $P(.)$, but then there is another probability distribution $P'(.)$ that would trouble strategy $\mathcal{S}'$.

The main problem is that knowledge of the proposed sequence of experiments permits a choice of distribution over hypotheses that is specifically geared towards making discovery hard for that sequence of experiments. A natural solution to this problem is to consider search strategies that do not commit to a particular experiment in light of a particular history of evidence, but rather to a distribution over possible experiments, i.e. a mixture of search strategies. This suggests a game-theoretic analysis.

4. Discovery as a Game

We can recast the above analysis of search strategies as a two person zero-sum game between Nature and the Scientist. The Scientist attempts to discover the true causal structure as efficiently as possible and Nature tries to make discovery as difficult as possible – in our case (for now) in terms of the number of experiments.

Nature initially gets to decide what the truth is – the underlying causal structure – but then has to stick with it, while the Scientist performs her experiments. Nature's pure strategies are all the directed acyclic causal structures over N variables. After each experiment by the Scientist the independence relations true in the underlying causal structure (manipulated by the intervention) are returned, i.e. the equivalence class of directed acyclic graphs that contains the true graph and is consistent with the sequence of experiments so far, is revealed. We refer to this – as is standard in game theory – as an information set. The pure strategies for the Scientist are all possible sequences of experiments. The Scientist may end the game after any sequence of experiments by declaring one of the graphs remaining in her information set as true. If the Scientist is correct, the payoff is the negative number of experiments that were performed (negative, since the Scientist wants to perform as few experiments as possible). If the Scientist is incorrect, the payoff is $-\infty$. Payoffs of $-\infty$ ensure that in order to avoid infinite loss the Scientist must be able to prove that her response is uniquely correct given the evidence.[4]

In game theory a strategy that specifies for each choice point a determinate choice (of experiment) corresponds to a *pure strategy*. A *mixed strategy* permits non-trivial distributions over the choices of experiments. Sometimes a mixed strategy can outperform any pure strategy. In our context the case for mixed strategies for Nature (i.e. distributions over graphs) is obvious – it would not be an interesting search problem if Nature were restricted to selecting one particular causal structure with probability 1. In the case of the Scientist we consider mixed strategies for two reasons. First, we already indicated at the end of Section 3 that there is no solution to the optimization problem when the Scientist is restricted to pure strategies. Second, and perhaps more intuitively, there are many circumstances in which a restriction to a specific experiment in light of the available evidence is artificial. For example, suppose there are two variables X and Y and it is known that either $X \to Y$ or $Y \to X$, each with probability 0.5. In that case a commitment to always intervene on X is artificial. Flipping a fair coin to either intervene in X or Y seems more appropriate.

For simplicity of exposition (and computation), we assume that every variable can be manipulated, that there are no latent variables and we only consider sequences of experiments in

4. Of course, one could integrate into the payoff structure some account of *how* wrong a Scientist is, but we leave this for future consideration.

which one (or no) variable is subject to an intervention per experiment.[5] Given that the worst case bound on the number of experiments under these circumstances is $N-1$ for $N>2$ variables, we do not need to consider search strategies for the Scientist that are longer than $N-1$ experiments, i.e. the table of results in Section 2 gives upper bounds on the worst case loss for the Scientist. Consider a simple example.

4.1 Example: Two Variables

Suppose there are just two variables. There are three possible causal structures among two variables X and Y, call them

$$\mathbf{Sa} := X \quad Y, \quad \mathbf{Sb} := X \rightarrow Y \quad \text{and} \quad \mathbf{Sc} := X \leftarrow Y.$$

Two experiments involving single interventions are sufficient and in the worst case necessary to discover the causal structure uniquely. The full game of Nature against Scientist is given in Figure 1. Nature can select among the three structures (grey boxes) **Sa**, **Sb** and **Sc**. The Scientist does not know which structure is selected, so **Sa**, **Sb** and **Sc** form an information set. The Scientist makes the next move and can end the game by guessing one of the structures without collecting any data (represented by the three arrows leaving each grey box upwards with Sa, Sb or Sc and the respective payoffs to Nature of 0 when the choice was correct and ∞ when incorrect). Alternatively, the Scientist can perform a passive observation (**N**), an intervention on the first variable (**X**), or an intervention on the second variable (**Y**). Depending on the choice and the true underlying graph, the game is either resolved because the graph can be uniquely identified (payoffs are indicated), or one of three new information sets – represented in the figure as a box containing the two causal structures that cannot be distinguished given the experiments so far – is returned. Again, the Scientist can end the game at this point with a guess, or can

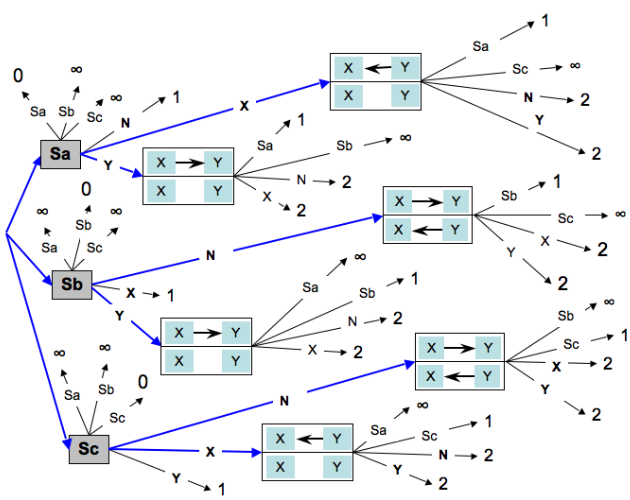

Figure 1: *Discovery of Causal Structure as a game of Nature against the Scientist, here for two causally sufficient variables.*

continue with a further experiment. Guesses and experiments that do not make sense in light

5. See Section 6 and Eberhardt (2007) for more on multiple simultaneous interventions.

of the evidence obtained so far, are not included in the game, since the Scientist is assumed to be rational. Given the worst case bound of two, there is no need to consider strategies of more than two experiments.

The game now permits an analysis of the optimal (mixed) search strategy against the most difficult probability distribution over causal structures. Since the game was constructed as a zero-sum game, the Nash equilibrium of the game corresponds to the mini-max solution, i.e. the Nash equilibrium specifies the desired upper bound on the expectation of the number of experiments sufficient and in the worst case necessary to discover the causal structure. The strategies implied by the Nash equilibrium for Nature and the Scientist, respectively, characterize a state, in which a unilateral change in strategy by Nature or by the Scientist does not improve their individual score.

An analysis of the game shows that the Nash equilibrium is given by a mixed strategy that is uniform over the three possible structures $\mathbf{Sa}, \mathbf{Sb}$ and $\mathbf{Sc}$ for Nature, and a mixed strategy for the Scientist that is uniform over passive observation, an intervention on X and an intervention on Y for the first experiment, and indifferent between possible (relevant) experiments for the second experiment, if a second experiment is necessary. That is, if Nature "selects" the true causal structure among the two variables uniformly, then Nature is making the discovery task maximally difficult for the Scientist. On the other side, by choosing uniformly whether to intervene on X, intervene on Y or just passively observe in the first experiment, the Scientist is doing the best she can to discover Nature's secrets efficiently, given that Nature is an adversarial player. Any other strategy, even mixed, will do no better and may well be worse (or will allow Nature to adapt accordingly to make things worse).

One could take this result to be a justification for the consideration of the uniform distribution over the hypothesis space in the assessment of an expected case performance for algorithms, but we will show below that this argument does not extend beyond the case of two variables.

The value of the Nash equilibrium represents the expected payoff to Nature (and loss to the Scientist) when playing the mixed strategy that is Nash. For this two variable game it is $5/3$ experiments, so the worst case expected performance is slightly better than the absolute worst case bound of 2 experiments. The Scientist's strategy is in this case not an equalizer, since some graphs are resolved in one experiment and others in two. The Nash equilibrium is, if we ignore the indifference for the second experiment, unique. As already indicated in the discussion of the optimization in the previous section, there is no Nash-equilibrium over pure strategies for either side, i.e. there is no Nash equilibrium if Nature selects one particular causal structure with probability 1, and there is no Nash equilibrium if the Scientist picks a experiments with probability 1. In both cases the opponent can adjust to do better. Further, returning with a guess of the true causal structure at any point is (obviously, given the infinities in the payoff structure) not Nash, so the solution that the Scientist returns is guaranteed to be justifiable given the evidence. Guessing (ending the game early) only becomes a viable option, when Nature is restricted to playing a subset of the possible structures.

The mixed strategy for the Scientist that is Nash is a Bayes solution, since it is a best response to the uniform distribution over structures. No two-experiment strategy (using single interventions) is a best response to any pure or mixed strategy by Nature. Interestingly, this last point does not apply in the case of three variable graphs. In the case of three variables, the game is substantially more complicated. There are 25 pure strategies for Nature (all DAGs over three variables) and 67 pure strategies for the Scientist (including all the early stops by guessing). We computed a Nash equilibrium, which determined 2 as the solution for the game: The worst case expected number of experiments sufficient and in the worst case necessary to determine the causal graph over three variables is two. That is, in the case of three variables, Nature can force the Scientist to the absolute worst case bound ($N-1=2$) even *in expectation*. To do so, Nature must select the true causal structure using a uniform distribution over the following set of 10 different graphs over three variables: the empty graph, three graphs consisting of a common effect only and the six possible complete graphs. Due to the edge-breaking nature of RCT-type interventions, at least two graphs of the 10 remain indistinguishable after any single intervention experiment or passive observation; hence a second experiment is necessary.

We know from the table in Section 2 that two experiments are sufficient for three variables. Consequently, the uniform distribution over the 10 graphs implies that any sequence of two different experiments is a best response, and obviously an equalizer (same payoff of two experiments, no matter which graph is true). No pure or mixed strategy will fare any better against the above distribution, which is not to say that there are no mixed strategies that do equally well.[6]

5. General Results

For single interventions per experiment, the three variable game is unique: For no other number of variables can Nature force the Scientist to the worst case bound in expectation. The general result for mixed strategies using single interventions per experiment is given by the following theorem.

Theorem *Given a set of $N > 3$ causally sufficient variables, the supremum of the expected number of experiments sufficient and in the worst case necessary to discover the causal structure is $\frac{2}{3}N - \frac{1}{3}$ experiments if only one (or no) variable can be subject to a RCT-type intervention per experiment.* ■

This bound is the value of a Nash equilibrium of the game: Nature plays a mixed strategy that is uniform over the *complete*(!) graphs over N variables only. For any N there are $N!$ such structures. From Nature's perspective, there is no advantage in considering incomplete causal structures, since for $N > 3$ variables, two single interventions have to be performed anyway, and in those two experiments any missing edge would be detected. This implies that a uniform distribution over all possible hypotheses (graphs) is not Nash for Nature, and an analysis based on such a distribution would underestimate the worst case expectation. For the Scientist the following strategy is Nash:

Strategy *Given N causally sufficient variables $X_1, \dots, X_n$, let each experiment $\mathcal{E}_i$ in the sequence intervene on $\mathbf{I}_i = \{X_j\}$, where X_j is selected uniformly from the variables that have not yet been subject to an intervention so far in the sequence.*

Since the game is symmetric with regard to the ordering of the variables (any variable can occur in any position in the graph), there are no order constraints on the Scientist's strategy. Of course, there may exist for some circumstances a particular order of experiments that minimizes the length of the sequence; but the Scientist cannot tell in advance.

For multiple simultaneous interventions per experiment the case is far more complicated. A discussion can be found in Eberhardt (2007); simulations suggest that the absolute worst case bound for multiple simultaneous interventions ($\lfloor \log_2(N) \rfloor + 1$ experiments) is fairly close to the worst case expectation, which would imply that the computationally simple pure strategies for the absolute worst case are fairly efficient even compared to the optimal mixed strategy, which is very hard to compute.

6. For example, if the passive observation is included as a possible first experiment, then if Nature chooses uniformly, it is a best response, but not an equalizer: $2/5$ of the time it finds the graph in one experiment and $3/5$ of the time it requires $8/3$ experiments (i.e. two experiments on average overall).

6. Conclusion

We framed the search for causal structure as a game in which Nature gets the first move to determine the graph after which the Scientist has free reign. This follows the approach developed for statistical hypothesis testing by Wald (1950), and generalizes it to sequences of experiments. Needless to say, this is only a first step presented with a very simple example. But the possibilities for generalization should now be obvious: (i) The effect of additional assumptions on the search procedure can be represented in terms of additional or reduced underdetermination in the information sets at any decision point. (ii) Cost other than the number of experiments can be considered. One may consider cost functions in terms of sample size, number of variables subject to intervention, or actual cost of experimentation – ethical or monetary. These cost functions need not be uniform across variables. (iii) Constraints or background knowledge on possible causal structures can be represented by limiting the possible pure strategies for Nature, while constraints on the set of experiments – e.g. it might not be possible to subject all variables to an experiment – limit the pure strategies for the Scientist. (iv) The robustness of search strategies can be analyzed in terms of changes in the optimal strategy with regard to off-equilibrium play by Nature – after all, Nature need not be adversarial; and the sensitivity of the optimal search strategy can be investigated by considering off-equilibrium play by the Scientist.

The game-theoretic approach to the discovery problem provides a general framework in which search strategies can be analyzed for their efficiency using a well-defined terminology and highly developed machinery. General guidelines for search procedures can be discovered and assessed on the basis of the explicit trade-off between discovery and its cost. Addressing these issues in the appropriate generality will require the integration of some of the most sophisticated game-theoretic techniques.

Acknowledgments

This research was supported by a fellowship from the James S. McDonnell Foundation. I am very grateful to Teddy Seidenfeld for making me think about this issue in the first place. Four anonymous reviewers (and one in particular, who was asked to review again) provided very useful comments on earlier drafts.

Appendix: Proofs[7]

Lemma 1 *For $N \geq 4$ the supremum of the expected number of experiments sufficient and in the worst case necessary to uniquely determine the causal graph is greater than 2 if only single interventions are permitted per experiment.*

Lemma 2 *The uniform distribution over complete graphs of N variables maximizes the expected number of experiments sufficient and in the worst case necessary to discover the true graph uniquely when only single intervention experiments are permitted.*

Theorem *Given a set of $N > 3$ causally sufficient variables, the supremum of the expected number of experiments sufficient and in the worst case necessary to discover the causal structure is $\frac{2}{3}N - \frac{1}{3}$ experiments if only one (or no) variable can be subject to a RCT-type intervention per experiment.* ■

Proof By Lemma 2, the uniform distribution over complete graphs is a worst case distribution. Suppose without loss of generality that the true complete graph over the variables $X_1, \ldots, X_N$ is such that for all $i < j$, $X_i \to X_j$. Under these circumstances an intervention on X_i is (1)

7. For more detailed proofs see Eberhardt (2007).

uninformative with respect to edge-orientation about all pairs of variables X_j, X_k with $j,k < i$; (2) uninformative with respect to edge-orientation about all pairs of variables X_j, X_k with $j,k > i$; and (3) informative for the remaining edges: It resolves (i) edges between variables X_j, X_k with $j > i > k$, (ii) outgoing edges from X_i and, (iii) since it is known that the graph is complete, edges broken by the intervention can be identified, and so all edges incident on X_i are resolved. In other words, an intervention on X_i splits the discovery problem into two subproblems, one with $N-i$ variables and the other with $i-1$ variables. About these subproblems, the intervention on X_i is uninformative.

Given the uniform distribution over complete graphs, the problem is entirely symmetric in the sense that each node is equally likely to be at any of the possible positions in a complete graph. Similarly, a uniform distribution selecting among the unintervened variables, implies that each variable is equally likely to be subject to an intervention in the first experiment. Consequently, we can give the expected number of experiments for this worst case distribution in terms of the numbers required for the subproblems the intervention creates:

$$E(\#\mathcal{E}(N)) = \frac{1}{N}\sum_{i=1}^{N}(E(\#\mathcal{E}(i-1)) + E(\#\mathcal{E}(N-i)) + 1) = 1 + \frac{2}{N}\sum_{i=1}^{N}E(\#\mathcal{E}(i-1))$$

where $E(\#\mathcal{E}(N))$ is the expected number of experiments required to discover the true graph if the graph is sampled from a Uniform over complete graphs of N variables. So the expected number of experiments for N variables is one plus the average of the sum of the number of experiments that it takes to resolve the two subproblems of size $N-i$ and $i-1$, respectively. For complete graphs with two and three variables, one can check by hand that $E(\#\mathcal{E}(2)) = 1$ and $E(\#\mathcal{E}(3)) = 5/3$. So finally we prove by induction that

$$E(\#\mathcal{E}(N)) = \frac{2}{3}N - \frac{1}{3} \qquad \text{for } N \geq 2.$$

It is true for $N = 2$. Suppose it is true for all integers up to some $N-1$. Then

$$E(\#\mathcal{E}(N)) = 1 + \frac{2}{N}\sum_{i=1}^{N}E(\#\mathcal{E}(i-1)) = 1 + \frac{2}{N}\sum_{i=1}^{N}(\frac{2}{3}(i-1) - \frac{1}{3}) = -\frac{1}{3} + \frac{2}{3}N$$

■

References

F. Eberhardt. *Causation and Intervention.* PhD thesis, Department of Philosophy, Carnegie Mellon University, 2007.

F. Eberhardt, C. Glymour, and R. Scheines. On the number of experiments sufficient and in the worst case necessary to identify all causal relations among n variables. In *Proceedings of the 21st Conference on Uncertainty and Artificial Intelligence*, pages 178–184, 2005.

Y. He and Z. Geng. Active learning of causal networks with intervention experiments and optimal designs. *Journal of Machine Learning Research*, 9:2523–2547, 2008.

S. Meganck, B. Manderick, and P. Leray. A decision theoretic approach to learning Bayesian networks. Technical report, Vrije Universiteit Brussels, 2005.

K. P. Murphy. Active learning of causal Bayes net structure. Technical report, Department of Computer Science, U.C. Berkeley, 2001.

S. Tong and D. Koller. Active learning for structure in Bayesian networks. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, pages 863–869. Morgan Kaufmann, 2001.

A. Wald. *Statistical Decision Functions*. Wiley, New York, 1950.

C. Yoo and G. Cooper. A computer-based microarray experiment design-system for gene-regulation pathway discovery. In *AMIA 2003 Symposium Proceedings*, pages 733–737, 2003.